

ABSTRACT

A reverse index useful for identifying documents in information retrieval searches may be used concurrently for indexing while it is updated with new documents.

Interruption to the use of the index is kept to a manageable level by partitioning the index and updating only single partitions of the index at a given time and further by bifurcating the index into a high speed supplemental portion that may be corrected concurrently on a real-time basis and which is periodically merged with the larger main portion. These two structures are merged during reading after brief locking, with pointer redirection.